APPARATUS, SYSTEM, AND METHOD TO PREVENT ADDRESS RESOLUTION CACHE SPOOFING

FIELD OF THE INVENTION

[0001] The present invention is generally related towards techniques to prevent spoofing of address resolution caches in computer systems.

BACKGROUND OF THE INVENTION

[0002] Individual computers of an Ethernet or other LAN network commonly have two addresses. The first address is an Internet Protocol (IP) address, which is a virtual network address assigned via a software source application. The IP address is used, for example, to form IP packets. However, each computer also has a physical address commonly known as the Medium Access Control (MAC) or also known as a hardware address. The MAC hardware source and destination addresses are necessary to prepare Ethernet headers to send data. Thus, in order for computers in an Ethernet network running IP to communicate, headers of individual frames require a source (sender) hardware address, and a destination (target) hardware address. This requires that a source computer station preparing to send a datagram to another station on the network know the correspondence between virtual IP addresses and physical hardware addresses.

[0003] There is a standard protocol known as the Address Resolution Protocol (ARP) that was initially designed to resolve Ethernet MAC addresses. Internet Request for Comments (RFC) 826 describes an ARP protocol for resolving Ethernet addresses, the contents of which are hereby incorporated by reference. RFC 826 describes a protocol to dynamically resolve correspondences between a network protocol address and a MAC address. ARP is not limited to operation on Ethernet; it is used to map IP addresses to MAC addresses on all types of broadcast-capable LAN networks. ARP includes a technique to request address resolution information, and a cache to temporarily save recently resolved MAC addresses.

[0004] Referring to prior art Figure 1, an ARP cache 100 comprises a table of correspondences between IP addresses 110 and MAC addresses 120. A local ARP cache 100 is maintained by a computer to map protocol addresses to hardware addresses. Conventionally, each time a computer receives an address resolution response it automatically updates ARP cache 100 with the sender's protocol address and hardware address. ARP cache 100 commonly has a finite size and is periodically flushed to eliminate obsolete entries.

[0005] Referring to prior art Figure 2, when a source application 205 in a computer station prepares a message having data 220 to be sent to a destination IP address 210 a lookup is performed by a MAC module 230 of the ARP cache 100 to determine a destination MAC address. If there is a cache hit for the destination MAC address 240, the destination MAC address is added to the frame that is sent out to the network.

[0006] However, if there is no cache entry in ARP cache 100 for the destination IP address, the source computer station broadcasts an ARP request message to the network. Referring to prior art Figure 3, the ARP message format includes fields for a sender's hardware address, sender's protocol address, target hardware address, target protocol address, hardware address type, protocol address type, hardware address length, protocol address length, and operation.

[0007] The broadcast ARP request message includes a source IP address, a source MAC address, and a target IP address. The broadcast ARP request message is a request for the computer that has the target IP address to respond back with its MAC address. The source computer then waits for a reply. The target station sends a unicast ARP reply to the source computer station with its IP address and its MAC address. The ARP cache 100 is updated and the source computer is now able to send a frame to the target.

[0008] However, there is a significant security issue associated with ARP. It is possible for ARP replies to be spoofed. Spoofing is a form of security breach in which a hacker masquerades as another user. In the context of LAN networks, spoofing includes inserting forged frames into the data stream. In ARP spoofing, a malicious entity creates forged ARP replies to corrupt the ARP cache with forged MAC addresses.

[0009] In one version of ARP spoofing, an ARP spoofer sends an unsolicited ARP reply with a spoofed MAC address for the IP address of a target computer. The recipient computer automatically updates its ARP cache 100 being updated with the spoofed MAC address. When the recipient computer tries to send data to the target computer, it ends up using the spoofed MAC address provided by the spoofer. This permits the spoofer to

intercept communications intended for the MAC address of the target computer. Additionally, ARP spoofing may be used to poison an ARP cache with erroneous MAC addresses so that data is lost.

[0010] ARP spoofing may also be used to initiate so-called "man in the middle" attacks, in which the spoofer creates spoofed MAC addresses in the ARP caches of a source computer and a destination computer which places the spoofer's computer in the middle of data flow between a source and target. Thus, if computer "A" wishes to send data to computer "B", a spoofer operating out of a computer "C" may place themself in the middle by creating a first spoofed MAC address in the ARP cache of computer A that fools computer A into sending data meant for computer B to computer C, and by creating a second spoofed MAC address in the ARP cache of computer B that fools computer B into sending data meant for computer A to computer C.

[0011] Other types of address resolution caches that are used to store an address resolution from a network protocol address to another type of address required to deliver data may be subject to similar types of spoofing. For example, IP Version 6 (IPv6) includes a neighbor discovery protocol. Address resolution in IPv6 is described in RFC 2461, the contents of which are hereby incorporated by reference. Address resolution in IPv6 includes the sending of multicast Neighbor Solicitation messages that include an IPv6 address of a target. A node having the IPv6 address responds with a Neighbor Advertisement indicating its IPv6 address and its link-layer address, where a link layer address is a link-layer identifier of an interface (e.g., IEEE 802 addresses for Ethernet and other LAN networks and E.164 addresses for ISDN networks). Additionally, a node may send an unsolicited Neighbor Advertisement to announce a link-layer address change.

[0012] In IPv6, a resolved link-layer address becomes an entry in a neighbor cache in the node. The link layer address resolution in IPv6 is thus analogous to ARP and the neighbor cache is analogous to the ARP cache. Consequently, the neighbor cache of IPv6 is potentially subject to analogous types of spoofing attacks in which a spoofer sends forged unsolicited Neighbor Advertisement messages with spoofed link-layer addresses.

Therefore, an improved apparatus, system, and method to prevent spoofing of an address resolution cache is desired.

SUMMARY OF THE INVENTION

[0013] An apparatus, system, method, and computer program product is disclosed for a firewall to prevent spoofing of an address resolution cache. An unsolicited message

received from a network that provides an address resolution for a network protocol address may provide either a genuine address resolution or be a spoofed message. In one embodiment, an unsolicited message that submits a first address resolution for a network protocol address is identified as a suspicious message if cached address resolution information has a second address resolution that differs from the first address resolution. Upon receiving a suspicious message, the accuracy of the first address resolution is checked. A request is issued for a network element having the specified network protocol address to reply with address resolution information. The suspicious message is determined to be a spoofed message if a reply is received that confirms that a network element claiming to own the specified network protocol address still has the first address resolution.

[0014] In one embodiment the firewall maintains a shadow copy of the address resolution cache that it uses to check cached address resolution information. In this embodiment, the firewall maintains the shadow copy of the address resolution cache and checks the shadow copy for address resolution information to determine if previously cached address resolution differs from a new address resolution. In some embodiments, the shadow copy has a greater residency lifetime than the original address resolution cache.

BRIEF DESCRIPTION OF THE FIGURES

[0015] The invention is more fully appreciated in connection with the following detailed description taken in conjunction with the accompanying drawings, in which:

[0016] Figure 1 illustrates a prior art Address Resolution Protocol (ARP) cache;

[0017] Figure 2 is a block diagram illustrating a prior art system for address resolution;

[0018] Figure 3 illustrates a prior art ARP message format;

[0019] Figure 4 is a block diagram of a system in accordance with one embodiment of the present invention;

[0020] Figure 5 is a flow chart of a method in accordance with one embodiment of the present invention;

[0021] Figure 6 is an interaction diagram showing a sequence of interactions for a normal address resolution update in accordance with one embodiment of the present invention; and

[0022] Figure 7 is an interaction diagram showing a sequence of interactions for an attempted spoofing of an address resolution cache in accordance with one embodiment of the present invention.

[0023] Like reference numerals refer to corresponding parts throughout the several views of the drawings.

DETAILED DESCRIPTION OF THE INVENTION

[0024] Figure 4 illustrates a network 400 in accordance with one embodiment of the present invention. A host computer station 405 is coupled to network 400 through a firewall 410. Network 400 includes at least one other computer station 409. In one embodiment, network 400 is a local area network (LAN). In one embodiment, each computer station 405 and 409 corresponds to an individual computer, such as a personal computer. However, it will be understood that an individual computer station 405 and 409 may correspond to a network element and that an individual network element may include a switch, router, server or other component for connecting network 400 to another network (not shown). Network 400 may also include one or more buses 402, switches, routers, or other network elements to couple data between the computer stations.

[0025] Host computer station 405 includes an address resolution cache 420. In one embodiment, address resolution cache 420 stores <network protocol address, data link layer address> pairs, where the network protocol address is a virtual address assigned by software and the data link layer address is an address required to deliver data to another network element. In one embodiment, address resolution cache 420 is a table associating network protocol addresses (e.g., IPv4 or IPv6 addresses) to data link layer addresses. Thus, each network protocol address in the cache has a corresponding data link layer address which is an address resolution for the network protocol address.

[0026] An address resolution module 417 uses the information in address resolution cache 420 in a lookup to determine a data link layer address for a network protocol address when it sends data to another computer station of network 400. For example, in an Address Resolution Protocol (ARP) embodiment, address resolution cache 420 may comprise an ARP cache having a table of IPv4 addresses and corresponding MAC addresses (i.e., <network protocol address, MAC address> pairs). For an IPv6 embodiment, address resolution cache 420 may be a neighbor discovery cache comprising a table of IPv6 addresses and corresponding MAC addresses. Address resolution cache 420 may be stored in a random access memory (RAM) or main memory of computer station 405.

[0027] Firewall 410 executes a protocol to prevent spoofing of address resolution cache 420 from a spoofer's computer station 407 that is directly or indirectly coupled to network 400. In one embodiment, an anti-spoofing state machine 430 within firewall 410

executes a protocol for identifying and preventing spoofing attacks directed at address resolution cache 420. As described below in more detail, firewall 410 examines cached address resolution information to identify suspicious address resolution messages that might be spoofed and then requests network elements to report address resolution information in order to determine whether the suspicious message is a real address resolution message or a spoofed address resolution message.

[0028] In one embodiment, firewall 410 includes a shadow copy 415 of an address resolution cache 420 of computer station 405 that it checks to identify messages that may originate from a spoofer. Shadow copy 415 includes the same type of address resolution stored in address resolution cache 420. However, since a conventional address resolution cache 420 commonly has a short cache residency lifetime before cache entries are evicted, in some embodiments of the present invention shadow copy 415 has an increased data size and residency lifetime compared to a conventional address resolution cache in order to improve the efficiency of firewall 410 at detecting spoofing attacks. However, throughout the following discussion it will be understood that in an alternate embodiment firewall 410 checks a conventional address resolution cache 420 that is modified to have an extended cache residency lifetime.

[0029] In one embodiment, firewall 410 comprises computer program instructions stored on a machine readable medium. For example, firewall 410 may comprise computer program instructions stored on a machine readable medium executable on computer station 405 or as computer program instructions stored on a machine readable medium which is executed on an intermediate hardware device disposed between computer station 405 and network 400. In an alternate embodiment, firewall 410 is disposed within a TCP/IP stack 417 associated with an operating system 419 of host computer station 405. Additionally, firewall 410 may be implemented on an intermediate device coupling host computer station 405 to network 400. It will thus be understood that firewall 410 may comprise a computer program product, be bundled with other computer software 419, be contained as part of a host computer station 405, or be contained within a hardware device coupling host computer station 405 to network 400. In one embodiment, firewall 410 is disposed on a Southbridge chipset, such as the nForce™ Media and Communications Processor (MCP) chipset manufactured by the Nvidia Corporation of Santa Clara, California.

[0030] Network 400 is a network in which MAC addresses are resolved by a protocol in which a computer station issues a request to resolve an address for a target network protocol address and a target computer station replies with address resolution information to

update an address resolution cache 420. Additionally, the network protocol provides for a computer to send an unsolicited reply to a computer station that includes address resolution information. In one embodiment, network 400 is a LAN network running IPv4 and address resolution cache 120 is an Address Resolution Protocol (ARP) cache. However, it will be understood throughout the following discussion that embodiments of the present invention may include other types of network protocols, such as Neighbor Discovery in the context of IPv6.

[0031] Figure 5 is a flow chart illustrating one embodiment of a method of preventing spoofing of an address resolution cache. Firewall 410 receives a message 505 submitting an unsolicited address resolution for a network protocol address (e.g., in an ARP embodiment an unsolicited ARP reply that includes the sender's IP address and MAC address). The message could be a genuine message sent by another (target) computer station 409 or a spoofed message sent by a spoofer's computer station 407. Consequently, firewall 410 checks cached addressed resolution information 510. In one embodiment firewall 410 checks the shadow copy 415 of address resolution information. (In an alternate embodiment, as previously discussed, if address resolution cache 420 has a sufficient residency lifetime, firewall 410 may check address resolution cache 420.)

[0032] If there is already a cache entry (a cache hit) having a different address resolution (e.g., a different MAC address in an ARP embodiment), firewall 410 requests 515 network elements (e.g., other computer stations) to report if they have an address resolution for the submitted network protocol address in order to check the accuracy/authenticity of the unsolicited new address resolution. In an ARP embodiment, this corresponds to sending a broadcast ARP request message. The computer station then waits a sufficient length of time to receive reply messages from other stations. If no reply message is received that matches the previously cached address resolution, firewall 410 determines that no attempt at spoofing has occurred and the address resolution cache 420 is updated 520 with the submitted address resolution. However, if a reply message includes the previously cached (old) address resolution, firewall 410 determines that an attempt at spoofing has occurred, and the previously cached address resolution is maintained 525 such that the cache is not updated with the submitted address resolution.

[0033] Figure 6 is an interaction diagram illustrating in more detail the situation where there is no attempt at spoofing and address resolution cache 420 is updated. An unsolicited address resolution message 605 is received from a target that has a genuine address resolution. The genuine address resolution may differ from a previously cached

address resolution for the IP address of the target. For example, in an ARP embodiment, a hardware upgrade may change the MAC address of the target computer station such that the target computer sends an unsolicited ARP reply to inform other network elements of its new MAC address. In the example of Figure 6, the firewall detects that the cached address resolution differs from the submitted address resolution. Since this is indicative of a suspicious message, the firewall acts to send a message 610 to network elements (e.g., other computer stations) for address resolution information for the network protocol address (e.g., an ARP broadcast message in an ARP embodiment). However, since in the example of Figure 6 no spoofing has occurred, only the target computer station (which is, in this case, apparently in sole possession of the target IP address) will respond with a reply message 615 that repeats the submitted address resolution for the target IP address. Consequently, since there is no contradiction between the reply message 615 and the submitted address resolution of the unsolicited address resolution message 605, the firewall may determine that the message is not spoofed and permit the submitted address resolution to be used to update the address resolution cache 420. Firewall 410 may, for example, permit address resolution cache 420 to be updated by permitting the genuine unsolicited message 605 or reply 615 to pass on to computer station 405.

[0034] Figure 7 is an interaction diagram illustrating in more detail the situation where spoofing is attempted. An unsolicited message 705 is received from a spoofer that has a spoofed address resolution for the network protocol address of a target computer station. The firewall checks for cached address resolution information. Since this example depicts an attempt at spoofing, the previous cache entry for the network protocol address contains the original address resolution for the target computer (e.g., the MAC address for the network protocol address of the target computer station). Consequently, the firewall determines that this is a suspicious message and issues a request message 710 to other network elements requesting address resolution information for the network protocol address to verify that the submitted address resolution is genuine. For this case, the target computer will reply 715 with the genuine address resolution, which is the previously cached entry. (Note that a spoofer will also typically reply 720 as well to such a broadcast message, repeating the spoofed address resolution).

[0035] In the example of Figure 7, the firewall can determine that an attempt at spoofing has occurred because the target computer issues a reply 715 that validates the original cache entry (and which disagrees with the spoofer's forged reply messages 705 and 720). As a result, firewall 410 prevents the address resolution cache from being updated.

The firewall 410 may block updating of the cache in a variety of ways, such as by blocking passage of replies 705 and 720 or by generating a signal to block updating address resolution cache 420.

[0036] In some embodiments of the present invention, firewall 410 also generates a report of spoofing attacks. The report may be stored on computer station 405 or sent to a system administrator.

[0037] One benefit of the present invention is that only a single computer station 405 needs a firewall 410 of the present invention to prevent address resolution spoofing of the computer station. Another benefit of the present invention is that it is compatible with widely used address resolution techniques, such as ARP, and does not require special hardware to be installed at each computer station of a network.

[0038] It will be understood that an embodiment of the present invention relates to a computer storage product with a computer-readable medium having computer code thereon for performing various computer-implemented operations. The media and computer code may be those specially designed and constructed for the purposes of the present invention, or they may be of the kind well known and available to those having skill in the computer software arts. Examples of computer-readable media include, but are not limited to: magnetic media such as hard disks, floppy disks, and magnetic tape; optical media such as CD-ROMs and holographic devices; magneto-optical media such as optical disks; and hardware devices that are specially configured to store and execute program code, such as application-specific integrated circuits ("ASICs"), programmable logic devices ("PLDs") and ROM and RAM devices. Examples of computer code include machine code, such as produced by a compiler, and files containing higher-level code that are executed by a computer using an interpreter. For example, an embodiment of the invention may be implemented using Java, C++, or other object-oriented programming language and development tools. Another embodiment of the invention may be implemented in hardwired circuitry in place of, or in combination with, machine-executable software instructions.

[0039] The foregoing description, for purposes of explanation, used specific nomenclature to provide a thorough understanding of the invention. However, it will be apparent to one skilled in the art that specific details are not required in order to practice the invention. Thus, the foregoing descriptions of specific embodiments of the invention are presented for purposes of illustration and description. They are not intended to be exhaustive or to limit the invention to the precise forms disclosed; obviously, many modifications and variations are possible in view of the above teachings. The embodiments were chosen and

described in order to best explain the principles of the invention and its practical applications, they thereby enable others skilled in the art to best utilize the invention and various embodiments with various modifications as are suited to the particular use contemplated. It is intended that the following claims and their equivalents define the scope of the invention.